

Vibrational Eigenstates of Four-Atom Molecules: A Parallel Strategy Employing the Implicitly Restarted Lanczos Method

R. B. Lehoucq,^a S. K. Gray,^b D.-H. Zhang,^c J. C. Light^c

^a *Mathematics and Computer Science Division, Argonne National Laboratory, Argonne, IL 60439. Current address: Sandia National Laboratories, MS 1110, P.O. Box 5800, Albuquerque, NM 87185-1110. E-mail: rlehoucq@cs.sandia.gov*

^b *Chemistry Division, Argonne National Laboratory, Argonne, IL 60439. E-mail: gray@tcg.anl.gov*

^c *Department of Chemistry, The University of Chicago, Chicago, IL 60636*

Abstract

We present an approach for determining the vibrational eigenstates of four-atom molecules. The primary representation of the (six-dimensional) eigenstates involves a finite basis or quantum number representation, whereas Hamiltonian matrix-vector products are evaluated with the aid of certain grid or discrete variable representations. This approach leads to computational and memory demands that are within acceptable limits. The implicitly restarted Lanczos method, as implemented in the ARPACK suites of codes, is then applied to determine some of the corresponding vibrational eigenstates. A distributed-memory parallel implementation of the method allows very large symmetric matrix eigenvalue problems—on the order of $N = 2 \cdot 10^6$ —to be tackled. The lowest fifty vibrational states of the HOCO molecule, with zero total angular momentum and even parity, are accurately computed on Argonne National Laboratory's IBM SP computer.

1 Introduction

The accurate determination of vibrational eigenstates is a challenging computational problem for many molecular systems [3,3]. The problem is particularly difficult when good zero-order pictures (e.g., normal modes) are not easily found or when highly excited states are desired. Two interrelated numerical tasks must be considered:

- determine a suitable finite representation of the underlying partial differential operator or Schrödinger equation, and

- determine the relevant eigenvalues and eigenvectors of the finite, but generally large symmetric Hamiltonian matrix \mathbf{H} that results.

When $N_a \geq 4$ atoms are involved and/or highly excited states are desired, the computation can easily suffer from *basis set explosion* because representations are typically product ones involving dimensions on the order of $N = n^F$, where n is on the order of 10 and F , the number of degrees of freedom, is $3N_a - 6$ (neglecting overall rotation). A representation that requires $n = 10$ to accurately represent the desired states then results in $N = 10^6$. Moreover, a poor choice of representation could result in n being significantly larger than 10, particularly if highly excited states are to be accurately described.

Once the representation is decided upon, the second task, determination of the eigenvalues and eigenvectors, may prove to be the computational bottleneck. For instance, the determination of *all* the eigenvalues when $N \geq 10^6$ is currently not feasible, nor is it desirable because only the lower energetic portion of the spectrum, or perhaps some particular spectral range, is typically of physical interest. Iterative methods such as the Lanczos method [3] can instead be used to determine some of the eigenvalues and eigenvectors. Since the Lanczos method requires matrix-vector products with \mathbf{H} , our representation must be such that an application of \mathbf{H} to a vector is computationally efficient.

Discrete variable representations (DVRs) [3,3,3] of the Schrödinger equation are useful because the resulting matrix-vector products are efficiently evaluated. The idea behind a one-dimensional DVR is to define a set of discrete points that is essentially equivalent to some more standard finite basis representation. The computational advantage arises from the fact that one part of the matrix-vector product, the “potential” part, is then well approximated by a diagonal matrix-vector product. When just one dimension is considered, there is no practical advantage of employing a DVR; but if two or more dimensions are being considered, the resulting matrix-vector product scales much better than N^2 . This improvement occurs because the potential part of the matrix-vector product remains diagonal and the kinetic part factors into blocks associated with the different dimensions. It is also possible to invoke energetic cut-offs within a DVR by excluding DVR points that yield potential energies greater than some cut-off value. This strategy reduces N and also the spectral range. Carrington and co-workers [3,3,3], for example, have developed and applied approaches to molecular vibrational states based on applying the Lanczos method to DVR or grid-based Hamiltonian representations with considerable success. Nevertheless, despite the advantages of such DVRs, or variations on the DVR idea such as potentially optimized [3,3] DVRs (PODVRs), the resulting product representations can still be very large.

Some applications of DVRs to vibrational eigenstate problems have employed the *sequential truncation-diagonalization* (STD) technique [3,3], which avoids

extremely large matrices by solving a sequence of reduced-dimensional eigenvalue problems. The technique fixes one or more of the coordinates in the problem at their DVR values, and represents the full eigenvalue problem in terms of the reduced-dimensional eigenvectors. (The use of a DVR representation makes it particularly convenient to carry out the STD idea.) The use of energetic cut-offs in deciding which reduced-dimension eigenvectors to employ can then lead to manageable effective values of N such that, for example, it may be possible to use standard eigenvalue-eigenvector methods. The STD approach is an excellent strategy that effectively reduces the dimension of the problem at the cost of additional algorithmic complexity.

The main objectives of the present paper are (i) to present an approach for obtaining vibrational eigenstates of four-atom systems using some of the ideas outlined above that, with the aid of parallel computers, can be applied to challenging problems and (ii) to give an illustration, with a nontrivial example, of how the approach works. In particular, we retain, in part, the idea of using DVRs or grid representations, but explore how a Lanczos method applied within a parallel computing environment may be used.

We employ an implicitly restarted Lanczos method (IRLM) as implemented in the parallel implementation [3] of the ARPACK [3] software package. The use of an implicitly restarted algorithm yields the corresponding eigenvectors directly. Such an approach can be more demanding in terms of computer memory than the straightforward Lanczos method. Therefore, to keep computer memory requirements down, we retain a quantum number (or finite basis) representation as the primary representation, and we employ the DVR or grid representations only as intermediate stages to facilitate the evaluation of the matrix-vector products.

Recently, Hayes and co-workers [3,3] have reported impressive calculations of the rovibrational eigenstates of certain three-atom systems using a combination of DVR/STD methods for the representation of the problems, and the IRLM for solving the resulting eigenvalue problems. A parallel implementation of the IRLM was shown to yield an efficient procedure for determining many eigenstates of a difficult molecular problem (HO_2 with total angular momentum $J > 0$). Our approach thus shares some similarities with this work. The main differences are that our focus is on larger, four-atom problems and we use a simpler combined quantum number/DVR Hamiltonian representation.

We illustrate the combined approach for a potential model of the HOCO molecule [3,3]. This system is challenging for at least two reasons. First, three relatively heavy atoms are present, leading to some relatively small de Broglie wavelengths. Second, both *cis* and *trans* isomers (with different effective normal mode frequencies) are important in the eigenstates we will determine. In fact, our calculations represent the best (and largest) theoretical vibrational

eigenstate estimates for HOCO currently available.

Section 2 describes our representation of the four-atom Hamiltonian matrix in terms of DVRs. Section 3 outlines the IRLM and our parallel implementation. Section 4 discusses the application to HOCO. Section 5 presents some brief concluding remarks.

2 Representation of the Hamiltonian Operator and Matrix-Vector Product

In this section we discuss both grid and quantum number representations.

2.1 Finite Basis and Grid Representations of the Hamiltonian Operator

We consider a system of four atoms A, B, C, and D interacting via a Born-Oppenheimer potential energy function V . While it is possible to consider nonzero total angular momenta J , we restrict attention to the case $J = 0$. In this case there are six internal degrees of freedom. Diatom-diatom Jacobi coordinates are employed. The three radial degrees of freedom are r_1 , the internuclear distance between A and B; r_2 , the internuclear distance between C and D; and R , the distance between the center of masses of AB and CD. The three angular degrees of freedom are θ_1 , the angle between the vectors associated with r_1 and R ; θ_2 , the corresponding angle associated r_2 and R ; and ϕ , the out-of-plane torsional angle [3,3]. The time-independent Schrödinger equation

$$\mathcal{H}\psi_k = \mathcal{E}_k\psi_k \quad , \quad k = 1, 2, \dots \quad , \quad (1)$$

for the vibrational eigenfunctions $\psi_k = \psi_k(R, r_1, r_2, \theta_1, \theta_2, \phi)$ and energies \mathcal{E}_k , involves the Hamiltonian operator (with $\hbar = 1$) [3,3]

$$\begin{aligned} \mathcal{H} = & -\frac{1}{2\mu} \frac{\partial^2}{\partial R^2} - \frac{1}{2m_1} \frac{\partial^2}{\partial r_1^2} - \frac{1}{2m_2} \frac{\partial^2}{\partial r_2^2} \\ & + \frac{\hat{j}_{12}^2}{2\mu R^2} + \frac{\hat{j}_1^2}{2m_1 r_1^2} + \frac{\hat{j}_2^2}{2m_2 r_2^2} + \mathcal{V}(R, r_1, r_2, \theta_1, \theta_2, \phi) \quad , \quad (2) \end{aligned}$$

with masses $\mu = (m_A + m_B)(m_C + m_D)/(m_A + m_B + m_C + m_D)$, $m_1 = m_A m_B / (m_A + m_B)$ and $m_2 = m_C m_D / (m_C + m_D)$. (The \hat{j}_1^2 , \hat{j}_2^2 , and \hat{j}_{12}^2 operators involve the angular variables and need not be explicitly defined here.)

The above Hamiltonian operator \mathcal{H} must be (approximately) represented in some finite matrix representation, \mathbf{H} . *Two* representations are employed: (i) a

quantum number (or finite basis) representation, \mathbf{H} , and (ii) a grid (or DVR) representation, \mathbf{H}^g . The quantum number representation, which is based on zero-order potentials adapted to the specific potential \mathcal{V} of interest, is our basic representation. For example, the distribution of the problem over processors and the iterative diagonalization procedure are all done within the quantum number representation. However, the actual (Hamiltonian) matrix-vector product, the key computational bottleneck, is evaluated using two representations, the quantum number representation for one part of the matrix-vector product, and the DVR representation for another part of it (due to the potential). This combined approach leads to the matrix-vector product scaling better than N^2 .

2.2 Quantum Number Representation

We may formally think of the quantum number representation in terms of abstract kets $|n, v_1, v_2, j_1, j_2, j_{12}\rangle$, with projection onto coordinate space

$$\langle R, r_1, r_2, \theta_1, \theta_2, \phi | n, v_1, v_2, j_1, j_2, j_{12} \rangle = \phi_n(R) \chi_{v_1}(r_1) \chi_{v_2}(r_2) F_{j_1, j_2, j_{12}}(\theta_1, \theta_2, \phi).$$

The basis functions for R , r_1 and r_2 are each based on suitable zero-order potentials, in other words are determined by the solution of the one-dimensional Schrödinger equations

$$\left[-\frac{1}{2\mu} \frac{\partial^2}{\partial R^2} + \mathcal{V}_R(R) \right] \phi_n(R) = \mathcal{E}_n \phi_n(R), \quad (3)$$

$$\left[-\frac{1}{2m_1} \frac{\partial^2}{\partial r_1^2} + \mathcal{V}_1(r_1) \right] \chi_{v_1}(r_1) = \mathcal{E}_{v_1} \chi_{v_1}(r_1), \quad (4)$$

$$\left[-\frac{1}{2m_2} \frac{\partial^2}{\partial r_2^2} + \mathcal{V}_2(r_2) \right] \chi_{v_2}(r_2) = \mathcal{E}_{v_2} \chi_{v_2}(r_2), \quad (5)$$

where \mathcal{V}_R , \mathcal{V}_1 , and \mathcal{V}_2 are defined, for example, by suitable one-dimensional cuts of the full potential \mathcal{V} . We denote the number of radial basis functions used for R , r_1 , and r_2 by N_R , N_1 , and N_2 , respectively.

The angular functions $F_{j_1, j_2, j_{12}}(\theta_1, \theta_2, \phi)$ are rotational eigenfunctions consistent with zero total angular momentum that are also chosen to be eigenfunctions of \hat{j}_1^2 , \hat{j}_2^2 , and parity. They are certain weighted sums of spherical harmonic functions associated with the two diatomic units. (Both j_1 and j_2 have allowed values $0, 1, 2, \dots$ and for $J = 0$ the allowed values of j_{12} are between $|j_1 - j_2|$ and $j_1 + j_2$. See Refs. [3,3] for further details.) We denote the total number of angular basis functions $F_{j_1, j_2, j_{12}}$ employed in our calculations as N_A .

Thus, in the quantum number representation a wavefunction is approximated as

$$\psi = \sum_{n,v_1,v_2,j_1,j_2,j_{12}} c_{n,v_1,v_2,j_1,j_2,j_{12}} \phi_n \chi_{v_1} \chi_{v_2} F_{j_1,j_2,j_{12}}, \quad (6)$$

where all sums are finite. The act of the Hamiltonian matrix in this quantum number representation on the vector \mathbf{c} is denoted $\mathbf{H}\mathbf{c}$.

2.3 Grid Representations

Let the kets $|R_i, r_{1j}, r_{2k}, \theta_{1a}, \theta_{2b}, \phi_c\rangle$ denote a formal grid or DVR, where $i, j, k, a, b,$ and c are discrete indices for the corresponding grid points or DVR. The $R_i, r_{1j},$ and r_{2j} points are taken to be those consistent with the associated basis function representations given above, following the ideas of Echave and Clary. Thus, the coordinate matrices associated with $R, r_1,$ and r_2 are diagonalized in the respective basis functions ($\phi_n, \chi_{v_1},$ and χ_{v_2}), with the eigenvalues corresponding to the desired (PO)DVR grid points. Since the PODVR points are adapted to the potential at hand, they lead to a particularly compact and accurate representation. The angular grid representation is a product one determined by Gauss-Legendre quadrature points for θ_1 and $\theta_2,$ and evenly spaced (Chebyshev or Fourier) points for $\phi.$ (The efficient representation of nondirect product basis functions such as spherical harmonics, and thus our angular functions, via angular grids is, of course, a difficult problem.)

If we represent $\mathbf{V}\mathbf{c},$ the potential part of the matrix-vector product in the quantum number representation, by employing the above-mentioned angular quadratures, we can then rewrite the matrix-vector product as a three-step process. First, transform the basis set to the grid representation; second, multiply by the diagonal potential matrix; third, transform back to the basis representation. The associated transformation matrices between the quantum number and angular grid representations thus involve the quadrature weights and the basis functions $F_{j_1,j_2,j_{12}}.$

An important reason for our use of an intermediate grid representation, as opposed to simply using just the quantum number or finite basis representation, lies in the angular part of the problem. (If one were to employ only a basis representation, there would be on the order of $N_1 N_2 N_R N_A^2$ associated angular potential matrix elements to store. As the total number of angular functions, $N_A,$ becomes large, e.g., on the order of 1000 or more, the number of matrix elements can exceed typical computer memory limitations. Of course, if one adopts an STD approach [3], this situation can be improved.)

Thus in the grid representation, the actual wavefunction is formally approxi-

mated as

$$\psi = \sum_{i,j,k,a,b,c} d_{i,j,k,a,b,c} f_i(R) f_j(r_1) f_k(r_2) g_{a,b,c}(\theta_1, \theta_2, \phi), \quad (7)$$

where f_i , f_j , f_k , and $g_{a,b,c}$ involve the underlying basis function representations and transformation functions.

Let M_A denote the total number of angular grid points, and let M_1 , M_2 , and M_R denote the number of radial PODVR grid points associated with r_1 , r_2 , and R . We note that better accuracy in the computed eigenvalues can be obtained by employing $M_i > N_i$, where N_i is the associated number of underlying basis functions in the quantum number representation. This is not surprising in relation to the angular degrees of freedom because a product grid representation is being employed to describe a nondirect product basis angular set involving spherical harmonics [3,3]. Furthermore, in practice (Sec. 4.1) we impose an energetic cut-off on the basis functions to be included that reduces N_A and further accentuates the difference between M_A and N_A .

Regarding the PODVRs used to represent the radial variables, it is also sometimes useful to employ $M_i > N_i$ because the associated quadratures over the potential are more accurately computed. (This is related to the fact that the underlying basis functions in a PODVR representation are not classical orthogonal polynomials, and so the procedure is not the same as employing a high-order Gaussian quadrature.) Allowing for $M_i > N_i$ leads to rectangular (as opposed to square transformation) matrices with more rows than columns but poses no real technical difficulties. The benefit is more accuracy for the same computer memory storage. This strategy also does not alter the spectral range of the Hamiltonian matrix, and so the number of iterations required to obtain the energy levels remains approximately the same. The Hamiltonian matrix remains a real symmetric matrix even when a rectangular transformation matrix \mathbf{Z} is involved.

2.4 Evaluating the Matrix-Vector Product

Our goal is the evaluation of the Hamiltonian matrix-vector product in the quantum number representation, $\mathbf{H}\mathbf{c}$. One can view the use of the grid representation (where the coefficients of the wavefunction are denoted by \mathbf{d}) as an intermediate step in obtaining $\mathbf{H}\mathbf{c}$.

The idea behind evaluating the matrix-vector product is most easily understood by a one-dimensional example. Consider a 1-D Hamiltonian operator $\hat{h} = \hat{t} + v(x)$, where \hat{t} is a differential kinetic energy operator and $v(x)$ a potential function. Let $|n_x\rangle$ be some quantum number representation corresponding

to the eigenstates of the zero-order Hamiltonian operator $\hat{h}_o = \hat{t} + v_o(x)$, that is, $\hat{h}_o|n_x\rangle = e_{n_x}|n_x\rangle$. Then, \mathbf{h} is the matrix defined by $\langle n'_x|h|n_x\rangle$, and the evaluation

$$\mathbf{h}\mathbf{c} = \mathbf{h}_o\mathbf{c} + (\mathbf{v} - \mathbf{v}_o)\mathbf{c} \quad (8)$$

follows. Since the quantum number representation is used for \hat{h}_o , it is a diagonal matrix with elements given by the corresponding eigenvalues e_{n_x} . Therefore, $\mathbf{h}_o\mathbf{c}$ is computationally fast. While $\mathbf{v} - \mathbf{v}_o$ is a full matrix in the quantum number representation, it may be taken to be a diagonal matrix in the grid representation based on \mathbf{h}_o . Thus, we introduce a DVR $|x_i\rangle$ which is related to the quantum number representation by a transformation matrix \mathbf{Z} such that $\mathbf{d} = \mathbf{Z}\mathbf{c}$ takes \mathbf{c} in the quantum number representation and yields the corresponding amplitudes \mathbf{d} in the grid representation. The algorithm for evaluating $\mathbf{h}\mathbf{c}$ is then

- (i) evaluate the act of the diagonal matrix \mathbf{h}_o on \mathbf{c} and save it as \mathbf{c}' ,
- (ii) transform from \mathbf{c} to \mathbf{d} via $\mathbf{d} = \mathbf{Z}\mathbf{c}$,
- (iii) evaluate the act of the diagonal matrix $\mathbf{g}^g - \mathbf{v}_o^g$ on \mathbf{d} via $\mathbf{d}'' = (\mathbf{v}^g - \mathbf{v}_o^g)\mathbf{d}$,
- (iv) transform from \mathbf{d}'' to \mathbf{c}'' using $\mathbf{Z}^T : \mathbf{c}'' = \mathbf{Z}^T\mathbf{d}''$, and
- (v) add \mathbf{c}'' to \mathbf{c}' , which yields the desired result.

Written together, the real symmetric nature of the effective potential portion of the Hamiltonian is

$$\mathbf{h}\mathbf{c} = \mathbf{h}_o\mathbf{c} + \mathbf{Z}^T(\mathbf{v}^g - \mathbf{v}_o^g)\mathbf{Z}\mathbf{c}. \quad (9)$$

If, for simplicity, we assume that N_x basis functions and grid points are involved, the resulting matrix-vector product still scales as N_x^2 , whether it is done in the above way or directly in the quantum number representation. Therefore, nothing is really gained by carrying out the matrix-vector product this way. However, when the problem is multidimensional, such as the present one, considerable numerical efficiency is gained. This occurs because the transformation matrices in the above procedure are direct products. For instance, if two coordinates are involved, say x and y , the numerical scaling of the matrix-vector product is proportional not to $(N_x N_y)^2$ but to $N_x N_y^2 + N_y N_x^2$.

The full six-dimensional matrix-vector product required for our problem is a straightforward generalization of the one- and two-dimensional discussion above. The quantum number representation is then defined by N_R basis functions associated with R , N_1 basis functions associated with r_1 , N_2 basis functions associated with r_2 , and N_A angular basis functions. The corresponding grid representations consist of M_R points in R , M_1 points in r_1 , M_2 points in r_2 , and M_A points in the three angular variables. Transformation matrices \mathbf{Z}_i (which are rectangular when $M_i > N_i$) associated with each degree of freedom

link basis and grid representations. The major computational bottlenecks in the matrix-vector product arise in applying the transformation matrices. We remark that the operation count of our approach is a complicated expression involving the N_i and M_i . However, the transformations between angular basis and quadrature points represent about 80% of the numerical effort.

3 Implicitly Restarted Lanczos Method and Parallel Implementation

We briefly review the Lanczos method, then discuss how to implicitly restart it, and finally how to implement a parallel version.

3.1 Implicitly Restarted Lanczos Method

The Lanczos method is a well-known method for calculating a few of the extremal eigenvalues of a large symmetric matrix. Given an initial starting (unit) vector q_1 , the method proceeds by computing the familiar three-term recurrence

$$\beta_j \mathbf{q}_{j+1} = \mathbf{H} \mathbf{q}_j - \alpha_j \mathbf{q}_j - \beta_{j-1} \mathbf{q}_{j-1}, \quad \text{where } \beta_0 \mathbf{q}_0 = 0 \quad \text{and } j \geq 1. \quad (10)$$

The eigenvalues of the symmetric tridiagonal matrix \mathbf{T}_j consisting of α_j 's and β_j 's on the diagonal and subdiagonal, respectively, are used to estimate those of \mathbf{H} . The extremal eigenvalues of \mathbf{T}_j quickly provide excellent estimates to those of \mathbf{H} . Since the method only requires knowledge of \mathbf{H} through its application to a vector, the method is well suited for large-scale eigenvalue problems.

Equation (10) may be rewritten in matrix form as

$$\mathbf{H} \mathbf{Q}_j = \mathbf{Q}_j \mathbf{T}_j + \beta_j \mathbf{q}_{j+1} \mathbf{e}_j^T, \quad (11)$$

where $\mathbf{Q}_j = \begin{bmatrix} \mathbf{q}_1 & \mathbf{q}_2 & \cdots & \mathbf{q}_j \end{bmatrix}$, \mathbf{T}_j is the associated tridiagonal matrix, and \mathbf{e}_j^T is the transpose of the last column of the order j identity matrix. We call this a Lanczos reduction of order j . The Lanczos vectors $\mathbf{q}_1, \dots, \mathbf{q}_j$ provide an orthogonal basis for the Krylov space $\{\mathbf{q}_1, \mathbf{H} \mathbf{q}_1, \dots, \mathbf{H}^{j-1} \mathbf{q}_1\}$, and thus \mathbf{T}_j is the orthogonal projection of \mathbf{H} onto the column span of this Krylov space. The Lanczos method is a generalization of the power method in that a sequence of iterates is used to approximate eigenvalues of \mathbf{H} .

Order the eigenvalues θ_i of \mathbf{T}_j so that $\theta_1 \leq \theta_2 \leq \cdots \leq \theta_j$, and denote the corresponding eigenvectors \mathbf{s}_i . Our goal is to present an algorithm for com-

putting k approximate eigenpairs $\theta_i, \mathbf{Q}_j \mathbf{s}_i (= \mathbf{y}_i)$ that satisfy the following two conditions:

- (i) $\|\mathbf{H}\mathbf{y}_i - \mathbf{y}_i\theta_i\| \leq \epsilon_U$ for $i = 1, \dots, k$ and
- (ii) $\mathbf{y}_i^T \mathbf{y}_j = \mathcal{O}(\epsilon_M)$ when $i \neq j$, and one otherwise.

The user-specified tolerance and machine precision are denoted by ϵ_U and ϵ_M , respectively. The first condition is that of accuracy, the second that of (numerical) orthogonality.

It is well known [3, pp. 479–485] that roundoff errors cause the Lanczos vectors to lose orthogonality. However, Paige [3] showed that the loss of orthogonality occurs precisely when an eigenvalue of \mathbf{T}_j is close to one of \mathbf{H} . In fact, the Lanczos vectors lose orthogonality in the direction of the associated approximate eigenvector. Moreover, failure to maintain orthogonality results in spurious copies of the approximate eigenvalue produced by the Lanczos method. We briefly review the traditional implementations of the Lanczos method that have emerged and their method for dealing with the loss of orthogonality.

- Implementations based on selective and partial orthogonalization [3,3]. These techniques carefully monitor the loss of orthogonality and perform additional orthogonalization steps only when necessary.
- A direct implementation of the three-term recurrence [3,3] with no additional orthogonalization steps.
- Orthogonalization of each new Lanczos vector against all the vectors generated.

The careful implementation of the first of these approaches is nontrivial. The second approach uses a strategy to determine whether the copies produced by the Lanczos method are spurious or correspond to a degenerate eigenvalue of \mathbf{H} . However, this approach may be used only for computing eigenvalues. Further computation is required if eigenvectors are of interest.

The third approach removes the complication associated with spurious copies, and eigenvectors are easily computed. Its drawbacks are the cost of maintaining orthogonality (on the order of Nm^2 floating-point operations for m Lanczos vectors) and storing the Lanczos basis vectors. If m is kept to a moderate size, the cost of maintaining the orthogonality of the Lanczos vectors is not a concern. We emphasize that the cost of maintaining full orthogonality of the Lanczos basis vectors may represent a minor cost (say, less than 5%) of the total cost in computing the eigenvalues and eigenvectors if m is not large and the cost of computing matrix-vector products $\mathbf{H}\mathbf{c}$ is large. Our numerical experiments will illustrate this behavior.

Suppose that we are able to compute m steps of (10) where m is chosen so that the cost of maintaining the orthogonality Lanczos vectors to machine

precision is small. Since we are interested in the k smallest eigenvalues of \mathbf{H} , consider starting another Lanczos method with some linear combination of $\mathbf{y}_1, \dots, \mathbf{y}_k$. This *restarting* of the Lanczos method is continued until the k smallest eigenvalues of \mathbf{T}_m satisfy the two aforementioned accuracy conditions. The motivation is that if $\mathbf{y}_1, \dots, \mathbf{y}_k$ are basis vectors for an eigenspace of \mathbf{H} , then $\beta_k = 0$ and the eigenvalues of \mathbf{T}_k are those of \mathbf{H} .

An equivalent manner in which to restart the Lanczos method with some linear combination of $\mathbf{y}_1, \dots, \mathbf{y}_k$ is by applying the matrix polynomial $(\mathbf{H} - \theta_{k+1}\mathbf{I}) \cdots (\mathbf{H} - \theta_m\mathbf{I})$ to the starting vector \mathbf{q}_1 . This removes the components from \mathbf{q}_1 corresponding to $\theta_{k+1}, \dots, \theta_m$ and specifies the linear combination of $\mathbf{y}_1, \dots, \mathbf{y}_k$. However, if matrix vector products with \mathbf{H} are expensive, then the application of this polynomial *filter* is expensive.

Sorensen [3] derived an elegant scheme for applying any polynomial filter of degree p ($\leq m$) *without* direct application of \mathbf{H} . Moreover, a length $m - p$ Lanczos reduction remains. We illustrate the scheme for a degree one polynomial. By subtracting $\mu\mathbf{Q}_m$ from a length m Lanczos reduction, we have

$$(\mathbf{H} - \mu\mathbf{I})\mathbf{Q}_m = \mathbf{Q}_m(\mathbf{T}_m - \mu\mathbf{I}) + \beta_m\mathbf{q}_{m+1}\mathbf{e}_m^T. \quad (12)$$

Now, we compute the orthogonal factorization $\mathbf{UR} = \mathbf{T}_m - \mu\mathbf{I}$, where \mathbf{U} is an orthogonal matrix and \mathbf{R} is upper triangular matrix. Substituting the orthogonal factorization into (12) results in

$$(\mathbf{H} - \mu\mathbf{I})\mathbf{Q}_m = \mathbf{Q}_m\mathbf{UR} + \beta_m\mathbf{q}_{m+1}\mathbf{e}_m^T. \quad (13)$$

A simple derivation gives

$$\mathbf{T}_m\mathbf{U} = \mathbf{U}(\mathbf{RU} + \mu\mathbf{I}) = \mathbf{UT}_m^+, \quad (14)$$

where \mathbf{T}_m^+ is an updated symmetric tridiagonal matrix. (Note that (14) is one step of the shifted QR algorithm [3, pp. 144–145].) Next, postmultiplying the original length m Lanczos reduction with \mathbf{U} results in

$$\mathbf{H}\mathbf{Q}_m\mathbf{U} = \mathbf{Q}_m\mathbf{T}_m\mathbf{U} + \beta_m\mathbf{q}_{m+1}\mathbf{e}_m^T\mathbf{U} = \mathbf{Q}_m\mathbf{UT}_m^+ + \beta_m\mathbf{q}_{m+1}\mathbf{e}_m^T\mathbf{U}, \quad (15)$$

where we used (14). Unfortunately (15) is no longer a length m Lanczos reduction. Because all the elements of \mathbf{U} below the subdiagonal band are zero ($\mathbf{T}_m - \mu\mathbf{I}$ is tridiagonal), dropping the last column of (15) gives a length $m - 1$ Lanczos reduction.

By equating the first $m - 1$ columns on both sides of (13), we see that the leading $m - 1$ columns of the updated Lanczos basis vectors $\mathbf{Q}_m\mathbf{U}$ provide an orthogonal basis for the same space as $(\mathbf{H} - \mu\mathbf{I})\mathbf{Q}_{m-1}$. (Note that the leading submatrix of order k in \mathbf{R} contains the orthogonalization (Gram–Schmidt) coefficients.) In particular, the updated starting vector has been replaced by

$(\mathbf{H} - \mu\mathbf{I})\mathbf{q}_1\rho^{-1}$, where ρ is the leading diagonal element of \mathbf{R} . Note that if $\mu = 0$, then a step of subspace iteration has been implicitly applied; this is analogous to the connection between the QR algorithm and subspace iteration. See [3] for further details.

The above scheme can be applied with each of the unwanted eigenvalues $\theta_{k+1}, \dots, \theta_m$. A length $m - p = k$ Lanczos reduction results where the k eigenvalues of \mathbf{T}_k are the wanted eigenvalues $\theta_1, \dots, \theta_k$ —the best k approximations to the smallest eigenvalues of \mathbf{H} . Equivalently, we have implicitly applied the matrix polynomial $(\mathbf{H} - \theta_{k+1}\mathbf{I}) \cdots (\mathbf{H} - \theta_m\mathbf{I})$ to the leading k Lanczos vectors.

Polynomial filters of degree no more than m may be applied in this fashion. Since the application of the polynomial filter is accomplished with an implicitly shifted QR algorithm on \mathbf{T}_m , the resulting algorithm is called the *implicitly restarted* Lanczos method (IRLM). In summary, the advantages of an IRLM include

- fixed storage requirements by fixing a value of m ,
- the ability to compute eigenvectors orthogonal to machine precision (since m is of modest size),
- application of the polynomial filter $\Psi(\mathbf{H})$ (of degree no larger than m) without the need for computing additional matrix vector products with \mathbf{H} , and
- incorporation of the well-understood numerical and theoretical behavior of the QR algorithm.

The last point is important because it allows the possibility of constructing general-purpose and reliable software for the large-scale eigenvalue problem. The reader is referred to [3] for further details.

3.2 Parallel Strategy

Suppose we have a length $j - 1$ Lanczos reduction $\mathbf{H}\mathbf{Q}_{j-1} = \mathbf{Q}_{j-1}\mathbf{T}_{j-1} + \beta_{j-1}\mathbf{q}_j\mathbf{e}_{j-1}^T$ and want to increase the length by one. The Lanczos matrix \mathbf{Q}_j is easily computed by appending \mathbf{q}_j to \mathbf{Q}_{j-1} , and the $j - 1$ subdiagonal element of \mathbf{T}_j is β_{j-1} . The computational burden is in computing the following three tasks:

- (i) $\mathbf{w} \leftarrow \mathbf{H}\mathbf{q}_j$,
- (ii) $\mathbf{g} \leftarrow \mathbf{Q}_j^T\mathbf{w}$,
- (iii) $\mathbf{f} \leftarrow \mathbf{w} - \mathbf{Q}_j\mathbf{g}$ (and thus $\beta_j = \|\mathbf{f}\|$ and $\mathbf{q}_{j+1} = \mathbf{f}/\beta_j$).

Steps (2) and (3) are performed so that the columns of \mathbf{Q}_j are orthogonal to \mathbf{q}_{j+1} . The vector \mathbf{f} is the projection of \mathbf{w} onto the orthogonal complement of the Krylov space $\{\mathbf{q}_1, \mathbf{H}\mathbf{q}_1, \dots, \mathbf{H}^{j-1}\mathbf{q}_1\}$. In practice, Steps (2) and (3) are

repeated when $\|\mathbf{f}\| < \sqrt{2}/2\|\mathbf{H}\mathbf{q}_j\|$; see [3,3] for further details. Note that \mathbf{g} is the j th column of \mathbf{T}_j ; β_{j-1} and α_j are the last two elements of \mathbf{g} , and the leading $j - 2$ elements are equal to zero.

The parallel implementation [3] of IRLM in ARPACK uses a data-parallel model of computation. So that large eigenvalue problems can be solved, the rows of the Lanczos matrix \mathbf{Q}_j are distributed (\mathbf{w} and \mathbf{f} are also conformably distributed) among the available processors. Because of the data distribution, parallelism is exploited during the computation of the matrix-vector product $\mathbf{H}\mathbf{q}_j$. Communication among the processors is required only for Step (2) and the computation of $\|\mathbf{f}\|$ and $\|\mathbf{w}\|$. Since each processor requires \mathbf{g} for Step (3), communication is accomplished with global reductions that are broadcast to every processor. Thus, every processor has a copy of \mathbf{T}_j and so performs the implicitly shifted QR algorithm needed for the implicit restarting.

Equation (6) suggests a straightforward data distribution over the radial variables in the quantum number (or finite basis) representation. Since \mathbf{H} has a simple block structure that operates on vectors of length N_A , the parallelization scheme spreads this work among the processors. (The application of the Hamiltonian matrix is distributed among at most $N_R N_1 N_2$ processors.) Our parallel strategy involves distributing consecutive rows of the Lanczos vectors in small multiples (at most 3) of size N_A among the available processors.

4 Vibrational States of HOCO

In this section we discuss the application of our model to HOCO.

4.1 Potential Model and Implementation Details

The four-atom system HOCO is of considerable theoretical and experimental interest because intermediate HOCO complexes are believed to play a significant role in the important combustion reaction $\text{OH} + \text{CO} \rightarrow \text{H} + \text{CO}_2$. See, for example, the discussion and references in [3,3]. Here we focus on determining some of the (bound) vibrational states of HOCO. This is a challenging theoretical problem owing to the importance of two energetically close isomers consistent with *trans* and *cis* conformations. Thus, normal mode descriptions welded to one particular equilibrium geometry are not appropriate, particularly for excited vibrational states. Furthermore, the presence of three relatively heavy atoms makes quantum mechanical studies of this system difficult, owing to the underlying small de Broglie wavelengths involved. A reasonable potential energy surface based on *ab initio* electronic structure data is readily

available [3,3]. (A newer surface [3] has been devised that improves some of the older surface’s energetics. Our main goal is to illustrate the utility of our computational method, and so the older surface [3,3] suffices for our purposes.)

In relation to the notation of Section 2, R now denotes the distance between the centres of mass of OH and CO; r_1 and r_2 denote the OH and CO internuclear distances; and θ_1 , θ_2 , and ϕ are the angular variables associated with OH and CO and the torsional angle, respectively. The radial DVR representations involve zero-order potentials based on the full potential in the region of the HOCO *trans* and *cis* isomers. The OH, CO and OH,CO zero-order potentials, used to define the radial grid representations, were chosen to be functions typical of the potential variation about the *cis* and *trans* equilibrium geometries. (Naturally there is no unique way of establishing such zero-order potentials. We simply fit Morse-like potential functions to the average of the potential variation in the *cis* and *trans* forms.)

We briefly point out a few features of the potential and related details. See [3,3] for a more thorough discussion. We employ an energy scale such that the potential energy is zero with OH + CO separated (large R limit) and with internuclear distances corresponding to minima in the relevant OH and CO parts of the potential. The potential energy minimum is a *trans*-HOCO structure with potential energy $-5.92 \cdot 10^{-2}$ a.u. (-1.61 eV or -12990 cm^{-1}). A second minimum, corresponding to *cis*-HOCO, exists at an energy of just $0.4 \cdot 10^{-2}$ a.u. (900 cm^{-1}) above the *trans* minimum, that is, at an energy of $-5.5 \cdot 10^{-2}$ a.u. with the energy zero as defined above. The barrier between *trans* and *cis* forms has potential energy $-4.74 \cdot 10^{-2}$ a.u. (Relative to the *trans* minimum, the barrier is $1.18 \cdot 10^{-2}$ a.u.)

The order of the Hamiltonian matrix in the quantum number representation is $N = N_R N_1 N_2 N_A$. Numerical experiments (see below) are carried out to test convergence of the computed energy levels with respect to the number of basis functions (N_R , N_1 , N_2) associated with the three radial degrees of freedom, and the number of basis functions (N_A) associated with the angular degrees of freedom. It is also necessary to test convergence with respect to the corresponding numbers of grid points M_i in the case of the radial variables where we allow $M_i \geq N_i$ for $i = R, 1$ and 2 . We note that N_A , in turn, is determined by the maximum number of j_1 and j_2 states employed, via standard angular momentum considerations [3,3]. Increasing the maximum allowed j_1 , j_1^{max} , or the maximum allowed j_2 , j_2^{max} , can significantly increase N_A .

We also employ certain energy cut-offs, which do not appreciably affect the computed eigenstates of interest. The potential energy is cut off at 0.1 a.u. (≈ 22000 cm^{-1}), which is well above the energies of interest and corresponds to the limit of unbounded OH + CO states. This cut-off significantly reduces

the spectral range (distance between maximum and minimum eigenvalues). Because the spectral range is modest (on the order of three orders of magnitude), the minimum (leftmost) eigenvalues of the discrete Hamiltonian may be computed with a Lanczos method that uses only matrix-vector products **Hc**.

It is also convenient to employ a zero-order rotational energy cut-off, which affects N_A , the total number of angular basis functions, by excluding any (j_1, j_2) combination that leads to the associated rigid rotor energy $B_1 j_1(j_1 + 1) + B_2 j_2(j_2 + 1)$, exceeding some value. Estimates B_1 and B_2 of the diatomic rigid rotor constants associated with OH and CO are used. (A rotational cut-off energy of 0.026 a.u., or 0.7 eV, was found to yield converged results, which we have checked by comparison with smaller energy cut-offs.)

4.2 Convergence of the First Fifty-Two Eigenstates

Six experiments were run to compute the lowest 52 vibrational states of the four-atom molecule HOCO. The largest matrix eigenvalue problem was of order 2,278,044.

For the largest eigenvalue problem, we set $N_A = 14,062$, with individual rotational quantum numbers up to $j_1^{max} = 28$ and $j_2^{max} = 54$. For the radial directions, $N_R = 9$, $N_1 = 6$ and $N_2 = 3$. The associated numbers of radial quadrature points were $M_R = 9$, $M_1 = 7$, and $M_2 = 5$, and the total number of angular quadrature points was $M_A = 46,255$.

This problem took approximately 10.7 hours using 54 nodes of the IBM SP system at Argonne National Laboratory. Since 54 processors were used, each processor stored 42,186 ($= 3 \cdot N_A$) rows of the Lanczos basis vectors. (Recall that the data distribution is done within the quantum number representation, and so up to $N_R N_1 N_2 = 162$ processors could be employed. However, given the configuration of the IBM SP, at most 60 processors were available.) ARPACK uses the MPICH [3] implementation of MPI [3] for communication among the processors.

Table 1 presents our best results for the lowest 52 energy levels ($J=0$, even parity) of HOCO with the assumed potential function in [3,3]. We emphasize that this study represents the first accurate determination of the vibrational eigenstates of HOCO. Some estimates of the first twenty vibrational energy levels have been published [3]. However, these were based on a 5D model that included only CO zero-point energy effects, and these levels, particularly beyond the first few levels, are not very accurate. Each of the first twenty energies in Table 1 is, in fact, lower than its corresponding counterpart in [3], with some of our higher energies lower by $7 \cdot 10^{-4}$ a.u. (150 cm^{-1}), which

Table 1
HOCO energy levels relative to the ground vibrational state $\mathcal{E}_1 = -8354.40 \text{ cm}^{-1}$.

1	0	14	1710.09	27	2242.55	40	2510.39
2	550.28	15	1726.35	28	2248.04	41	2537.76
3	553.19	16	1777.18	29	2273.97	42	2540.66
4	1044.23	17	1850.25	30	2312.55	43	2546.72
5	1092.03	18	1855.85	31	2314.01	44	2596.97
6	1098.84	19	1950.20	32	2360.05	45	2613.81
7	1180.14	20	1960.82	33	2362.17	46	2639.62
8	1218.11	21	2030.75	34	2370.88	47	2642.36
9	1329.64	22	2057.10	35	2382.47	48	2674.12
10	1520.28	23	2076.02	36	2388.86	49	2676.64
11	1557.32	24	2110.35	37	2391.59	50	2732.11
12	1595.25	25	2146.34	38	2414.37	51	2744.82
13	1618.92	26	2212.85	39	2472.09	52	2768.75

is consistent with our more accurate calculations. (A detailed comparison of the results of the two calculations is beyond the scope of the present paper.) We have inspected some of the corresponding eigenstates, but have not made a detailed study. (The inspection of such multidimensional eigenstates is a difficult matter because one must consider a variety of cuts and averages over the six-dimensional space.) The ground vibrational state is, as expected, a *trans*-HOCO structure. It is consistent with a zero-point energy of $\approx 2.1 \cdot 10^{-2}$ a.u. (4630 cm^{-1}). However, the excited states may be *cis* or *trans*-like or perhaps exhibit features consistent of the two isomers. This is because of the small $0.33 \cdot 10^{-2}$ a.u. (726 cm^{-1}) energy difference between *cis* and *trans* isomers, and a relatively low barrier connecting the two isomers. We plan to present a more detailed discussion of the eigenstates in future work.

We also mention that of the remaining five experiments, three differed only in the number of angular functions used from our large problem. The three experiments set N_A equal to 8, 481, 10, 115 and 12, 434. The order of the corresponding discrete Hamiltonian were 1, 373, 922, 1, 638, 630 and 2, 014, 308, respectively. The total time needed to compute the 52 eigenpairs was 6.2, 7.6, and 8.9 hours. This shows that our parallelization scheme scales linearly with the number of angular basis functions. For all these experiments, the lowest vibrational levels agreed to five significant digits, and the highest levels agreed to three significant digits. The remaining two experiments used discrete

Hamiltonians of order 1,000,000.

4.3 Computational Details Associated with IRLM in P_ARPACK

A few remarks are in order regarding the use of the parallel implementation of IRLM. For all six experiments, the IRLM maintained the numerical orthogonality of all the Lanczos vectors to machine precision (approximately 15 decimal digits). This was accomplished by reorthogonalizing at every step of the Lanczos reduction whenever necessary.

The size m at which the Lanczos reduction was restarted never exceeded a value of 82. The default strategy of the ARPACK implementation of IRLM applied the minimum of $(m - k)/2$ and $m - k - s$ unwanted eigenvalues per restart where s denotes the number of eigenvalues that satisfy the accuracy condition (k is the number of requested eigenvalues). This strategy results in better performance than applying $m - k$ unwanted eigenvalues as shifts at every restart.

We used a value of $\epsilon_U = 10^{-7}$ so that $\|\mathbf{H}\mathbf{y}_i - \mathbf{y}_i\theta_i\| \leq \epsilon_U$. We remark that only the last (or, equivalently, the largest) 10% of the wanted eigenpairs satisfied the accuracy condition. The remaining ones satisfied the condition to a tolerance of machine precision. This situation occurs because the initial eigenpairs are further refined as the last eigenpair finally satisfies the accuracy condition. For all size experiments, an average of 17 matrix-vector products per eigenvalue was needed to satisfy the accuracy condition. The eigenvectors \mathbf{y}_i were also orthogonal to machine precision.

We emphasize that the cost in time to maintain the full orthogonality of the Lanczos basis vectors along with all the other costs associated with the parallel ARPACK implementation of IRLM—excluding the matrix-vector products—represented only 2%–3% of the total computation time. Thus, the cost of the computation is completely dominated by the cost of a matrix-vector product $\mathbf{H}\mathbf{c}$.

We also remark that roughly half of the time needed to complete a matrix-vector product involved communication time. A better communication strategy may be undertaken with a resulting increase in performance per processor and overall performance. For example, after the initial data distribution in the quantum number representation (6), a further distribution in the grid representation (7) should be undertaken. (At present, we do not parallelize over the M_R grid points because $M_i > N_i$ for $i = R, 1, 2$ leads to a more complicated implementation.)

This improved communication strategy would also remove redundant work

arrays (used on every processor) so that larger problems can be solved for a fixed number of processors. Nonetheless, our goal was to demonstrate that extremely large eigenvalue problems can be solved on distributed memory machines while avoiding out-of-core memory methods. A further data distribution in the grid representation would lead to larger eigenvalue problems solved more efficiently.

5 Concluding Remarks

We outlined an approach for obtaining vibrational energy states of four-atom systems. Our approach utilizes grid and discrete variable representation ideas and a parallel implementation of the IRLM for the calculation of eigenvalues and eigenvectors. We considered a very challenging problem, the HOCO molecule, and obtained over fifty accurate eigenstates. Our results, the first accurate determination of the vibrational levels for this system, should also provide a useful benchmark for testing other theoretical methods for eigenvalue determination of large systems. While we did not discuss in detail the physical nature of the eigenstates, we note that one of the strengths of the current procedure is that it yields directly (without need of subsequent calculations) the eigenstates as well as the eigenvalues.

The main purpose of our particular parallel strategy was to make possible the determination of many vibrational eigenstates of the HOCO problem via the IRLM. As pointed out in the preceding section, communication between processors, with our current parallel strategy, occupies a significant portion of the matrix-vector product. Therefore, we did not investigate in detail any performance issues such as the scaling of the method with respect to the number of processors used. Better strategies can be adopted to improve communication time, and these will be investigated in future work.

We also plan to examine the physical nature of the HOCO eigenstates, extend our results for this system, and consider other interesting four-atom systems. Our current approach focuses on energy states from the ground state up, and, because of the memory requirements, it is difficult to obtain, say, hundreds of eigenstates for very large problems such as the HOCO example. Thus, as with the general numerical eigenvalue problem, an important future direction is to devise and implement suitable methods for determining eigenstates in specific energy ranges. For example, the approach outlined here could be applied with a shift-and-invert transformation or Green function techniques [3]. Filter-diagonalization methods [3,3] also appear to be interesting alternative approaches to the problem.

Acknowledgements

RBL was supported by the U.S. Department of Energy, Contracts DE-FG0f-91ER25103 and W-31-109-Eng-38. SKG was supported by the Office of Basic Energy Sciences, Division of Chemical Sciences, U.S. Department of Energy under Contract No. W-31-109-ENG-38. JCL and SKG are grateful for a University of Chicago-Argonne National Laboratory collaborative grant. JCL also gratefully acknowledges use of the IBM SP computer in the Center for Computational Science and Technology at Argonne National Laboratory.

References

- [1] Z. Bacic and J. C. Light. Theoretical methods for rovibrational states of floppy molecules. *Annu. Rev. Phys. Chem.*, 40:469–498, 1989.
- [2] E. L. Sibert. Variational and perturbative descriptions of highly excited molecules. *Int. Rev. in Phys. Chem.*, 9:1–27, 1990.
- [3] C. Lanczos. An iteration method for the solution of the eigenvalue problem of linear differential and integral operators. *J. Research of the National Bureau of Standards*, 45:255–282, 1950.
- [4] J. V. Lill, G. A. Parker, and J. C. Light. Discrete variable representations and sudden models in quantum scattering theory. *Chem. Phys. Lett.*, 89:483, 1982.
- [5] J. C. Light and Z. Bacic. Adiabatic approximation and nonadiabatic corrections in the discrete variable representation: Highly excited vibrational states of triatomic molecules. *J. Chem. Phys.*, 87:4008, 1987.
- [6] J. C. Light. Discrete variable representations in quantum mechanics. In J. Broeckhove and L. Lathouwers, editors, *Time-Dependent Quantum Molecular Dynamics: Experiments and Theory*, pages 185–199. Plenum, New York, 1992.
- [7] P. N. Roy and T. Carrington. A direct-operation Lanczos approach for calculating energy levels. *Chem. Phys. Lett.*, 257:98–104, 1996.
- [8] M. J. Bramley and T. Carrington. Calculation of triatomic vibrational eigenstates - product or contracted basis sets, Lanczos or conventional eigensolvers - what is the most efficient combination. *J. Chem. Phys.*, 101:8494–8507, 1994.
- [9] M. J. Bramley and T. Carrington. A general discrete variable method to calculate vibrational energy levels of 3-atom and 4-atom molecules. *J. Chem. Phys.*, 99:8519–8541, 1993.
- [10] J. Echave and D. C. Clary. Potential optimized discrete variable representation. *Chem. Phys. Lett.*, 190:225–230, 1992.

- [11] H. Wei and Jr. T. Carrington. The discrete variable representation of a Triatomic Hamiltonian. *J. Chem. Phys.*, 97:3029–3037, 1992.
- [12] Z. Bacic and J.C. Light. Accurate localized and delocalized vibrational states of HCN/HNC. *J. Chem. Phys.*, 86:3065–3077, 1987.
- [13] K. J. Maschhoff and D. C. Sorensen. P_ARPACK: An efficient portable large scale eigenvalue package for distributed memory parallel architectures. In Jerzy Wasniewski, Jack Dongarra, Kaj Madsen, and Dorte Olesen, editors, *Applied Parallel Computing in Industrial Problems and Optimization*, volume 1184 of *Lecture Notes in Computer Science*, Berlin, 1996. Springer–Verlag.
- [14] R. B. Lehoucq, D. C. Sorensen, and C. Yang. *ARPACK USERS GUIDE: Solution of Large Scale Eigenvalue Problems by Implicitly Restarted Arnoldi Methods*. SIAM, Philadelphia, PA, To appear in 1998. Also available as technical report CRPC-TR96665, Center for Research on Parallel Computation, Rice University, Houston, TX.
- [15] X.T. Wu and E. F. Hayes. Ho₂ rovibrational eigenvalue studies for non-zero angular momentum. *J. Chem. Phys.*, 107:2705–2718, 1997.
- [16] P. P. Korambath, X.T. Wu, and E. F. Hayes. Enhanced method for determining rovibrational eigenstates of van der Waals molecules. *J. Phys. Chem.*, 100:6116–6124, 1996.
- [17] G. C. Schatz, M. S. Fitzcharles, and L. B. Harding. State-to-state chemistry with fast hydrogen atoms: Reaction and collisional excitation in H + CO₂. *Faraday Discuss. Chem. Soc.*, 84:359–369, 1987.
- [18] K. Kudla, G. C. Schatz, and A. F. Wagner. A quasiclassical trajectory study of the OH + CO reaction. *J. Chem. Phys.*, 95:1635–1647, 1991.
- [19] D. H. Zhang and J. Z. H. Zhang. Photofragmentation of HF dimer: Quantum studies on *ab initio* potential energy surfaces. *J. Chem. Phys.*, 99:6624–6633, 1993.
- [20] D. H. Zhang and J. Z. H. Zhang. Quantum calculations of reaction probabilities for HO + CO → H + CO₂ and bound states of HOCO. *J. Chem. Phys.*, 103:6512–6519, 1995.
- [21] G. C. Corey and D. Lemoine. Pseudospectral method for solving the time-dependent Schroedinger equation in spherical coordinates. *J. Chem. Phys.*, 97:4115–4126, 1992.
- [22] J. Dai and J. C. Light. A grid representation for spherical angles. *J. Chem. Phys.*, in press, 1997.
- [23] G. H. Golub and C. F. Van Loan. *Matrix Computations*. Johns Hopkins University Press, Baltimore, third edition, 1996.
- [24] C. C. Paige. *The computation of eigenvalues and eigenvectors of very large sparse matrices*. PhD thesis, University of London, London, 1971.

- [25] B. N. Parlett. *The Symmetric Eigenvalue Problem*. Prentice-Hall, Englewood Cliffs, N.J., 1980.
- [26] H. Simon. Analysis of the symmetric Lanczos algorithm with reorthogonalization methods. *Linear Algebra and Its Applications*, 61:101–131, 1984.
- [27] J. Cullum and R. A. Willoughby. Computing eigenvalues of very large symmetric matrices—an implementation of a Lanczos algorithm with no reorthogonalization. *J. Comp. Phys.*, 434:329–358, 1981.
- [28] B. N. Parlett and J. Reid. Tracking the progress of the Lanczos algorithm for large symmetric eigenproblems. *IMA Journal of Numerical Analysis*, 1:135–155, 1981.
- [29] D. C. Sorensen. Implicit application of polynomial filters in a k-step Arnoldi method. *SIAM J. Matrix Analysis and Applications*, 13:357–385, 1992.
- [30] R. B. Lehoucq. Truncated QR algorithms and the numerical solution of large scale eigenvalue problems. Preprint MCS-P648-0297, Mathematics and Computer Science Division, Argonne National Laboratory, Argonne, Ill, 1997.
- [31] E. M. Goldfield, S. K. Gray, and G. C. Schatz. Quantum dynamics of a planar model for the complex forming $\text{OH} + \text{CO} \rightarrow \text{H} + \text{CO}_2$ reaction. *J. Chem. Phys.*, 102:8807–8817, 1995.
- [32] K. S. Bradley and G. C. Schatz. A quasiclassical trajectory study of $\text{H} + \text{CO}_2$: Angular and translational distributions, and OH angular momentum alignment. *J. Chem. Phys.*, 106:8464–8472, 1997.
- [33] W. Gropp, E. Lusk, N. Doss, and A. Skjellum. A high-performance, portable implementation of the MPI message passing interface standard. *Parallel Computing*, 22:789–828, 1996.
- [34] Message Passing Interface Forum. MPI: A message-passing interface standard. *International Journal of Supercomputing Applications*, 8:165–414, 1994.
- [35] R. E. Wyatt. Matrix spectroscopy: Computation of interior eigenstates of large matrices using layered iteration. *Phys. Rev. E*, 51:51–3658, 1995.
- [36] M. R. Wall and D. Neuhauser. Extraction, through filter-diagonalization, of general quantum eigenvalues or classical normal mode frequencies from a small number of residues or a short-time sequence of a signal. I. Theory and application to a quantum dynamics model. *J. Chem. Phys.*, 102:8011–8022, 1995.
- [37] V. A. Mandelshtam and H. S. Taylor. A low-storage filter diagonalization method for quantum eigenenergy calculation or for spectral analysis of time signals. *J. Chem. Phys.*, 106:5085–5090, 1997.